Recommendations for using Multiple IP Addresses in Benchmarking Tests

draft-lencse-bmwg-multiple-ip-addresses

Gábor LENCSE lencse@sze.hu (Széchenyi István University) – presenter
Keiichi SHIMA keiichi.shima@g.softbank.co.jp (SoftBank)
Ole TROAN ot@cisco.com (Cisco)

IETF 123, BMWG, July 21, 2025.

Outline

- The progress of the draft
- Reminder
  - Problem description: why testing with multiple IP addresses is needed?
  - Recommended Solution: Usage of multiple, pseudorandom IP addresses, and the test setup with gateways
- Request for WG adoption

Progress of the Draft

- “03” Presented at IETF 121 (in person)
  - New measurement setups using gateways
  - Investigation of the number of IP addresses to be used
  - Results of benchmarking tests with and without using gateways were presented
  - Recommendation for testing according to the setup using gateways
- “04” Addressed the comments of Boris Hassanov (not presented)
- “05” current version at IETF 123 (online)
  - Addressed the comments from Giuseppe Fioccola and Libin Liu
  - Ole Troan joined us as a co-author

Reminder: Problem Description: Conditions

- RFC 2544 has defined a test frame format with fixed IP addresses and fixed port numbers.
- RFC 4814 introduced pseudorandom port numbers, but it kept the usage of a single source and destination IP address pair when a single destination network is used.
- Receive Side Scaling (RSS) supports the receiving of multi-million packets per second by distributing the load among the CPU cores
  - Depending on implementation, the hash function includes:

    1st type: source IP address, destination IP address

    2nd type: source IP address, destination IP address, source port, destination port

Reminder: Problem Description: Unfairness

- RFC 4814 pseudorandom port numbers + 2nd RSS implementation
  - Works perfectly (port numbers ensure entropy)
  - All CPU cores are used, load is distributed approximately evenly
- RFC 4814 pseudorandom port numbers + 1st RSS implementation
  - Gives poor results (no entropy is ensured as IP addresses are fixed)
  - Thus only two CPU cores are used (one core per direction)
- However, network interconnect devices using the 1st RSS implementation work perfectly, when they forward Internet traffic (IP addresses ensure entropy)

→ Conditions for the laboratory tests should be improved!

Reminder: Recommended Solution

- Basic idea: Let us use pseudorandom IP addresses!
  - This is the spirit of RFC 4814 applied to the IP addresses ☺
- Problems to solve:
  - What ranges can be used?
    - There is scarcity in IPv4 addresses reserved for benchmarking
      - 198.18.0.0/15 was reserved for benchmarking
    - There is abundance in IPv6 addresses reserved for benchmarking
      - 2001:2::/48 was reserved for benchmarking
  - What ranges should be used?
    - A trade-off was pointed out
    - And a better solution was proposed ☺

Reminder: When Gateways Are Used (IPv4)

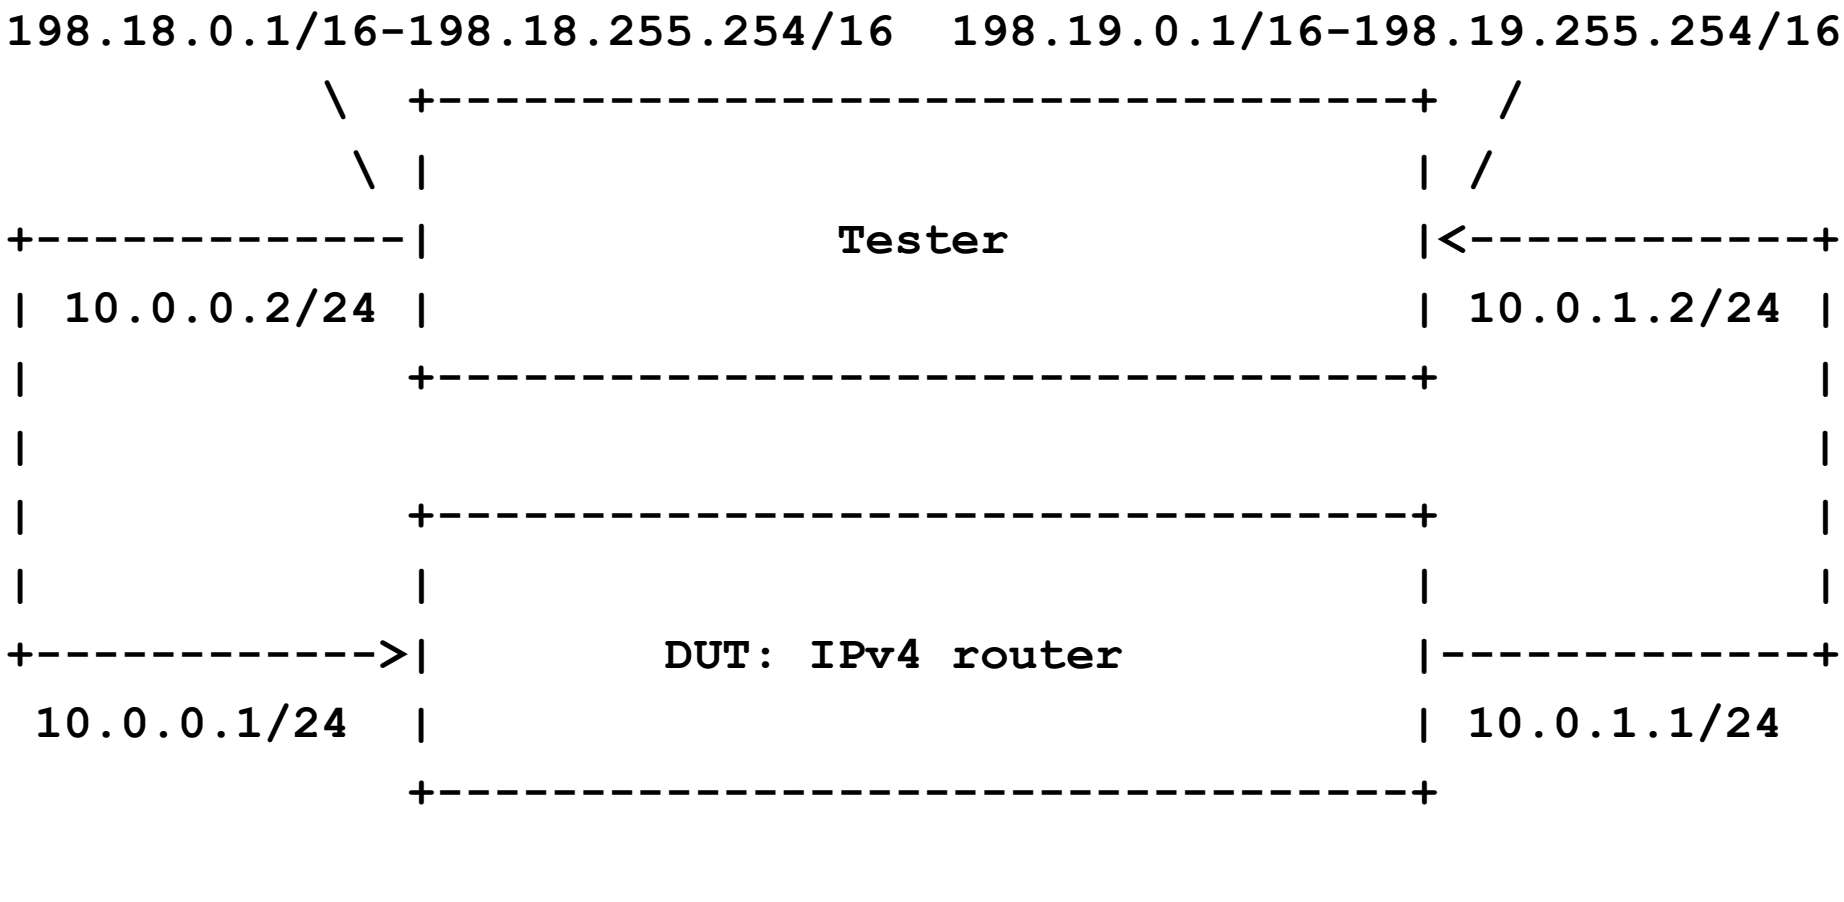

routes in the DUT:

destination network	next hop
198.18.0.0/16	10.0.0.2
198.19.0.0/16	10.0.1.2

We proposed

- The usage of multiple IP addresses
  - To support the 1st type of RSS
- The usage of a gateway
  - To avoid the overhead caused by the high number of ARP / ND table entries
- The usage of all available IP addresses (in their ranges)
  - 64k-2 for IPv4 and 64k for IPv6

We would like to ask for WG adoption

We believe that our method is usable in practice and it significantly improves the conditions of laboratory tests by removing their bias against the 1st type of RSS implementation.

- Would you support the adoption of the draft as a WG item?